arXiv:1401.3424v3 [math.ST] 1 Dec 2014

Marčenko-Pastur Law for Tyler's and Maronna's M-estimators

Teng Zhang
The Program in Applied and Computational Mathematics (PACM), Princeton University
Princeton, New Jersey 08544, USA
tengz@princeton.edu

Xiuyuan Cheng
Département d'informatique, École normale supérieure, 45 rue d'Ulm, Paris 75005, France
xiuyuan.cheng@ens.fr

Amit Singer
Department of Mathematics and PACM, Princeton University
Princeton, New Jersey 08544, USA
amits@math.princeton.edu

December 2, 2014

Abstract

This paper studies the limiting behavior of Tyler's and Maronna's M-estimators, in the regime that the number of samples n and the dimension p both go to infinity, and p/n converges to a constant y with $0 < y < 1$. We prove that when the data samples are identically and independently generated from the Gaussian distribution $N(\mathbf{0}, \mathbf{I})$, the difference between $\frac{1}{n}\sum_{i=1}^{n} \boldsymbol{x}_i \boldsymbol{x}_i^T$ and a scaled version of Tyler's M-estimator or Maronna's M-estimator tends to zero in spectral norm, and the empirical spectral densities of both estimators converge to the Marčenko-Pastur distribution. We also prove that when the data samples are generated from an elliptical distribution, the limiting distribution of Tyler's M-estimator converges to a Marčenko-Pastur-Type distribution.

1 Introduction

Many statistical estimators and signal processing algorithms require the estimation of the covariance matrix of the data samples. When the underlying distribution of the input data samples $\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_n \in \mathbb{R}^p$ is assumed to have zero mean, a commonly used estimator is the sample covariance matrix $\boldsymbol{S}_n = \frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i \boldsymbol{x}_i^T$.

However, the estimator $\boldsymbol{S}_n$ is sensitive to outliers, and performs poorly in terms of statistical efficiency (i.e., it has a large variance) for heavy-tailed distributions, whose probability densities decay slower than the Gaussian density as $\|\boldsymbol{x}\| \to \infty$. For these cases, a commonly used robust estimator of covariance is Maronna's M-estimator [17] $\bar{\Sigma}$, which is defined as the solution to

$$\bar{\Sigma} = \frac{1}{n}\sum_{i=1}^n u(\boldsymbol{x}_i^T \bar{\Sigma}^{-1} \boldsymbol{x}_i) \boldsymbol{x}_i \boldsymbol{x}_i^T, \tag{1}$$

where $u : (0, \infty) \to [0, \infty)$ is continuous, nonincreasing and $xu(x)$ is strictly increasing. These assumptions on $u(x)$ are necessary for the existence and uniqueness of the solution to (1) [16, Theorem 2.2].

Another interesting robust covariance estimator is Tyler's M-estimator [23] $\hat{\Sigma}$, which is the unique solution to

$$\hat{\Sigma} = \frac{p}{n}\sum_{i=1}^n \frac{\boldsymbol{x}_i \boldsymbol{x}_i^T}{\boldsymbol{x}_i^T \hat{\Sigma}^{-1} \boldsymbol{x}_i}, \ \operatorname{tr}(\hat{\Sigma}) = 1. \tag{2}$$

Tyler's M-estimator can be viewed as a special case of Maronna's M-estimator with $u(x) = \frac{p}{x}$, although this particular $u(x)$ does not guarantee the uniqueness and existence of the solution to (1). In fact, without the assumption $\operatorname{tr}(\hat{\Sigma}) = 1$, the solution to (2) would not be unique, since for any $c > 0$, $c\hat{\Sigma}$ is also the solution to (2).

Tyler's M-estimator gives the "shape" of the covariance, but is missing its magnitude. However, we remark that for many applications the "shape" of the covariance suffices, for example, the principal components can be obtained from $\hat{\Sigma}$.

Compared with the sample covariance estimator, Tyler's M-estimator is more robust to heavy-tailed elliptical distributions. The density function of elliptical distributions takes the form

$$f(\boldsymbol{x}; \Sigma, \mu) = |\Sigma|^{-1/2} g\{(\boldsymbol{x} - \mu)^T \Sigma^{-1} (\boldsymbol{x} - \mu)\},$$

where g is some nonnegative function such that $\int_0^\infty x^{p-1} g(x)\, \mathrm{d}x$ is finite. This family of distributions is a natural generalization of the Gaussian distribution by allowing heavier/lighter tails while maintaining the elliptical geometry of the equidensity contours. Elliptical distributions are considered important in portfolio theory and financial data, and we refer to [11, Section 4] for further discussion. Besides, elliptical distributions are used in the radar data [21], where the empirical distributions are heavy-tailed because of outliers.

It has been shown that when a data set follows from an unknown elliptical distribution (with mean zero), Tyler's M-estimator is the most robust covariance estimator in the sense of minimizing the maximum asymptotic variance [23]. This property suggests that Tyler's M-estimator should be more accurate than the sample covariance estimator for elliptically distributed data sets. Also empirically, it has been shown to outperform the sample covariance estimator in applications such as finance [13], anomaly detection in wireless sensor networks [6], antenna array processing [20] and radar detection [21].

1.1 Asymptotic analysis in a high-dimensional setting

Many scientific domains customarily deal with sets of large dimensional data samples, and therefore it is increasingly common to work with data sets where the number of variables, p, is of the same order of magnitude as the number of observations, n. Under this high-dimensional setting, the asymptotic spectral properties of $\boldsymbol{S}_n$ at the limit of infinite number of samples and infinite dimensions have been well studied. A noticeable example is the case that the entries of $\{\boldsymbol{x}_i\}_{i=1}^n$ are independent identically distributed random variables with mean 0 and variance 1. Denoting the eigenvalues of $\boldsymbol{S}_n$ by $\lambda_1(\boldsymbol{S}_n), \lambda_2(\boldsymbol{S}_n), \cdots, \lambda_n(\boldsymbol{S}_n)$, the Marčenko-Pastur law [18] states that when $p, n \to \infty$ and $p/n \to y$, where $0 \leq y \leq 1$, the distribution of the eigenvalues of $\boldsymbol{S}_n$, i.e. the empirical spectral density

$$f_{n,p}(\lambda) = \frac{1}{p} \sum_{i=1}^{p} \delta_{\lambda_i(\boldsymbol{S}_n)}(\lambda)$$

converges in distribution to the Marčenko-Pastur distribution defined by

$$\rho_{\mathrm{MP},y}(x) = \frac{1}{2\pi} \frac{y\sqrt{(y_+ - x)(x - y_-)}}{x} \mathbf{1}_{[y_-, y_+]}, \text{ where } y_\pm = (1 \pm \sqrt{y})^2. \quad (3)$$

The limiting empirical spectral density of Maronna's M-estimator when both $p, n \to \infty$ and $p/n \to y$ has been analyzed in two recent works [7, 8], where it is shown that a properly scaled Maronna's M-estimator converges to $\boldsymbol{S}_n$ in operator norm under the assumptions that $u(x)$ is nonnegative, nonincreasing and continuous; $xu(x)$ is nondecreasing and bounded and $\sup_x xu(x) > 1$. To the best of our knowledge, no similar result has been obtained for Tyler's M-estimator, although it has been conjectured that the empirical density distribution follows the Marčenko-Pastur distribution [13, 8]. Some works focused on the case $p, n \to \infty$ and $p/n \to 0$: Dümbgen [10] showed that the condition number of Tyler's estimator converges to $1 + O(\sqrt{p/n})$, and Frahm and Glombek [12] showed that the empirical spectral distribution of $\sqrt{n/p}(\bar{\Sigma} - \mathbf{I})$ converges to a semicircle distribution.

1.2 Spiked covariance model

A particular application of the high-dimensional analysis of the covariance estimators is the "spiked covariance model" [15]. This model arises naturally from

applications where the clean signal comes from a low-dimensional "signal subspace" and the noise is isotropic. As $p, n \to \infty$, one considers a sequence of covariance estimation problems that satisfy the following two assumptions:

- The number of observations n and the number of variables p in the n-th problem grow proportionally: $p/n \to y$, $0 < y < 1$.
- The observations are sampled from a distribution with zero mean and covariance Σ_p, and the eigenvalues of Σ_p are given by $(l_1, l_2, \cdots, l_r, 1, 1, \cdots, 1)$, where the number of spikes r and the amplitudes $l_1 \geq l_2 \geq \cdots \geq l_r \geq 1$ do not depend on n.

A fundamental problem in the spiked covariance model is that, given a set of observations, what is number of "spikes" in the estimated covariance? That is, what is r, and in particular, is r nonzero? It is natural to expect that the eigenvalues of the sample covariance matrix play a basic role in answering this question. Indeed, the distribution of the eigenvalues of the sample covariance matrix have been well studied in the "null" case, the case with no spike and $\Sigma_p = \mathbf{I}$. Here are some several known results:

- Empirical spectral densities [18]: for any real x, the empirical density distribution of the sample covariance converges to $\rho_{\mathrm{MP},y}(x)$.
- The largest eigenvalue converges to $\rho_{\mathrm{MP},y}(y_+)$ almost surely, and the distribution of the largest eigenvalues (appropriately scaled) approaches the Tracy-Widom distribution [15].
- The robustness to models: The Tracy-Widom distribution holds for a quite general class of independent, identically distributed random samples [22].

The nonnull case has also been extensively studied in [3, 2, 19].

Empirical evidence motivates the theoretical study of the spike covariance model through Maronna's and Tyler's M-estimators instead of the sample covariance estimator: It is shown in [23] that for elliptically distributed data sets, Tyler's M-estimator might perform better than the sample covariance estimator as shown in. Moreover, it has been empirically shown that for heavy-tailed data sets, Maronna's M-estimator performs better than the sample covariance matrix for the multiple signal classifier (MUSIC) algorithm [Section III][7].

To embark on this theoretical study, one needs to generalize the existing theory of the sample covariance matrices to Maronna's and Tyler's M-estimators. In this paper we prove the convergence of the empirical spectral density and largest eigenvalue of these estimators in Section 3.3 and Corollary 3.5, and partially provide theoretical guarantees for their performance in the "null" case. However, we remark that this paper is only a first step towards a complete understanding of Maronna's and Tyler's M-estimators in spiked covariance models: we do not specify the distribution of the largest eigenvalue in the null case, and the nonnull case has not been studied yet.

1.3 Main results

In this paper, we analyze Tyler's and Maronna's M-estimators in the high-dimensional setting. Our main results, Theorem 3.6 and Corollary 3.8, show that as $p, n \to \infty$ and $p/n \to y$, $0 < y < 1$, the empirical spectral density of a properly scaled Tyler's M-estimator converges to the Marčenko-Pastur distribution $\rho_{\mathrm{MP},y}(x)$. Based on the properties of Tyler's and Maronna's M-estimators, this paper also analyzes the empirical spectral density when data samples are i.i.d. drawn from other distributions, such as elliptical distributions.

When data samples are generated from elliptical distributions, the empirical spectral density of the sample covariance estimator has been studied in [11, Theorem 2]. Compared to Corollary 3.10, the limiting spectral distribution of $\boldsymbol{S}_n$ is much more complicated, and therefore our result might be more applicable in practice.

Both analyses in this paper and in [7, 8] are based on the representation of the M-estimator as a weighted sum of $\boldsymbol{x}_i\boldsymbol{x}_i^T$, and the uniform convergence of the corresponding weights. However, we give a different proof for the convergence of the weights, by considering the weights as the solution to an optimization problem or a system of equations. In comparison, our approach can handle Tyler's M-estimator and some Maronna's M-estimators (i.e., some $u(x)$) that are not covered by the results in [7, 8]. We remark that while some Lemmas and technical proofs also appear in [7, 8] (for example, Lemma 5.2 and the analysis in the proof of Theorem 3.2 are similar to [7, Lemma 2], [8, Lemma 6] and the proof of Theorem 1 in [7]), we still include them for the completeness of the paper.

The rest of the paper is organized as follows. In Section 2 we provide the definition of Tyler's and Maronna's M-estimators and state some of their properties, such as existence and uniqueness, and we introduce their representations as linear combinations of $\boldsymbol{x}_i\boldsymbol{x}_i^T$. In Section 3 we present the main results that when the data set is i.i.d. sampled from the Gaussian distribution $N(\mathbf{0}, \mathbf{I})$, properly scaled Tyler's and Maronna's M-estimators converge to $\boldsymbol{S}_n$ in operator norm, and the limiting empirical spectral density of Tyler's M-estimator follows the Marčenko-Pastur law. We also extend the result to elliptical distributions for Tyler's M-estimator and non-isotropic Gaussian distributions for Maronna's M-estimator. The technical proofs are given in Section 5.

As for notations, we will use c, c', C, C' to denote any fixed constants as $p, n \to \infty$ (though they may depend on y). Depending on the context, they might denote different values in different equations.

2 Properties of Tyler's and Maronna's M-estimators

When $\mathrm{span}(\{\boldsymbol{x}_i\}_{i=1}^n) = \mathbb{R}^p$, Tyler's M-estimator exists and is unique [26, Theorem 1.1].

As for Maronna's M-estimator, for the convenience of high-dimensional analysis we define it slightly different from the literature by removing the factor $1/n$

from (1), or equivalently, replace $\frac{1}{n}u(x)$ by $u(x)$ in (1):

$$\bar{\Sigma} = \sum_{i=1}^{n} u(\boldsymbol{x}_i^T \bar{\Sigma}^{-1} \boldsymbol{x}_i) \boldsymbol{x}_i \boldsymbol{x}_i^T, \tag{4}$$

and we note that similar modifications are applied in other works on high-dimensional analysis of Maronna's M-estimator [7, 8].

The existence and uniqueness of Maronna's M-estimator has been analyzed in [16, 27], by analyzing the minimizer of the objective function

$$L(\Sigma) = \sum_{i=1}^{n} \rho(\boldsymbol{x}_i^T \Sigma^{-1} \boldsymbol{x}_i) + \frac{n}{2} \log\det \Sigma, \text{ where } \rho'(x) = nu(x)/2.$$

The derivative of $L(\Sigma)$ with respect to Σ^{-1} is

$$\frac{\mathrm{d}}{\mathrm{d}\Sigma^{-1}} L(\Sigma) = \sum_{i=1}^{n} \frac{n}{2} u(\boldsymbol{x}_i^T \Sigma^{-1} \boldsymbol{x}_i) \boldsymbol{x}_i \boldsymbol{x}_i^T - \frac{n}{2}\Sigma,$$

and therefore the stationary points of $L(\Sigma)$ are the solutions to (4).

By the geodesic convexity of $L(\Sigma)$, [27, Theorem 1] states that the uniqueness of the minimizer of $L(\Sigma)$ is guaranteed when $\rho(x)$ is continuous in $(0, \infty)$, nondecreasing and $\rho(e^x)$ is convex, and the minimizer of $L(\Sigma)$ exists when $a_1 = \sup\{a | x^{a/2} \exp(-\rho(x)) \to 0 \text{ as } x \to \infty\}$ is positive [16, Theorem 2.3] (when $\lim_{x\to\infty} xu(x)$ exists, $a_1 = n \lim_{x\to\infty} xu(x)$),[1] and

$$\frac{|\{\boldsymbol{x}_i\}_{i=1}^n \cap \mathrm{V}|}{n} < 1 - \frac{p - \dim(\mathrm{V})}{a_1} \text{ for any linear subspace } \mathrm{V} \in \mathbb{R}^p. \tag{5}$$

When the above condition holds, the minimizer of $L(\Sigma)$ exists, and the minimizer is also a solution to (4). Due to the geodesic convexity of $L(\Sigma)$ [27, Theorem 1], the minimizer of $L(\Sigma)$ is unique, and any solution to (4) is also a minimizer of $L(\Sigma)$. Therefore, the solution to (4) is also unique. That is, under the above assumptions on $\rho(x)$, we have the existence and uniqueness of the solution to (4).

Since uniqueness and existence of both Maronna's M-estimator and Tyler's M-estimator requires $\mathrm{span}(\{\boldsymbol{x}_i\}_{i=1}^n) = \mathbb{R}^p$, the case $p > n$ or $y > 1$ does not apply to these estimators and throughout the paper we assume $y < 1$.

The analysis for Tyler's and Maronna's M-estimators in this paper is based on the following representations, whose proofs are deferred to Section 5. We remark that equation (7) in Lemma 2.1 has appeared in [25, (27)] and [14, Section A] as "covariance estimation in scaled Gaussian distributions" and "Barthe's convex program", but its connection to Tyler's M-estimator has not been rigorously justified.

[1] It follows from the comment after [16, Definition 2.1]. We remark that $u(x)$ in [16] should be replaced by $nu(x)$, since we use (4) over the standard definition (1). This also explains our choice of $\rho'(x) = nu/2$ instead of $\rho'(x) = u/2$ used in [16]. We remark that there is a typo after [16, Definition 2.1], where "$\rho'(x) = 2u(x)$" should be replaced by "$\rho'(x) = u(x)/2$".

Lemma 2.1. *Tyler's M-estimator can be written as*

$$\hat{\Sigma} = \sum_{i=1}^{n} \hat{w}_i \boldsymbol{x}_i \boldsymbol{x}_i^T \Big/ \mathrm{tr}\Big(\sum_{i=1}^{n} \hat{w}_i \boldsymbol{x}_i \boldsymbol{x}_i^T \Big), \tag{6}$$

where $\{\hat{w}_i\}_{i=1}^n$ *are uniquely defined by*

$$(\hat{w}_1, \hat{w}_2, \cdots, \hat{w}_n) = \underset{w_i>0, \sum_{i=1}^n w_i = 1}{\arg\min} - \sum_{i=1}^{n} \log w_i + \frac{n}{p} \log\det\big(\sum_{i=1}^{n} w_i \boldsymbol{x}_i \boldsymbol{x}_i^T \big). \tag{7}$$

Lemma 2.2. *When Maronna's M-estimator exists and is unique, any* $\{\bar{w}_i\}_{i=1}^n$ *satisfying*

$$\bar{w}_j = \boldsymbol{x}_j^T \Big(\sum_{i=1}^{n} u(\bar{w}_i) \boldsymbol{x}_i \boldsymbol{x}_i^T \Big)^{-1} \boldsymbol{x}_j, \ \textit{for } j = 1, 2, \cdots, n \tag{8}$$

gives Maronna's M-estimator by

$$\bar{\Sigma} = \sum_{i=1}^{n} u(\bar{w}_i) \boldsymbol{x}_i \boldsymbol{x}_i^T \, . \tag{9}$$

3 Main Results

In this section we present the main results: we prove the convergence of Tyler's and Maronna's M-estimators to $\boldsymbol{S}_n$ under the Gaussian model $N(\mathbf{0}, \mathbf{I})$ in terms of the operator norm in Section 3.1, and then extend the result to elliptical distributions/non-isotropic Gaussian distributions in Section 3.2. Based on the convergence, we obtain the limiting empirical density distributions of Tyler's and Maronna's M-estimators in Section 3.3.

3.1 Isotropic Gaussian Distribution

3.1.1 Tyler's M-estimator

In this section, we assume that $\{\boldsymbol{x}_i\}_{i=1}^n \subset \mathbb{R}^p$ are i.i.d. drawn from $N(\mathbf{0}, \mathbf{I})$. The main result, Theorem 3.2, characterizes the convergence and convergence rate of Tyler's M-estimator to $\boldsymbol{S}_n$ in terms of the operator norm. Its proof applies Lemma 3.1, whose proof is rather technical and therefore deferred to Section 5.

Lemma 3.1. *If* $\{\boldsymbol{x}_i\}_{i=1}^n$ *are i.i.d. sampled from* $N(\mathbf{0}, \mathbf{I})$*, then* $\max_{1\le i\le n} |n\,\hat{w}_i - 1|$ *converges to* 0 *almost surely as* $p, n \to \infty$*. In particular, there exist* $C, c, c' > 0$ *such that for any* $\varepsilon < c'$*,*

$$\Pr\left(\max_{1\le i\le n} |n\,\hat{w}_i - 1| \le \varepsilon \right) \ge 1 - Cne^{-c\varepsilon^2 n}. \tag{10}$$

Theorem 3.2. *Suppose that $\{\boldsymbol{x}_i\}_{i=1}^n$ are i.i.d. sampled from $N(\mathbf{0},\mathbf{I})$, $p,n\to\infty$ and $p/n=y$, where $0<y<1$, and $\boldsymbol{x}_i\sim N(\mathbf{0},\mathbf{I})$ for all $1\le i\le n$, then a scaled Tyler's M-estimator converges to $\boldsymbol{S}_n$ in operator norm almost surely, and there exist $C,c,c'>0$ such that for any $\varepsilon<c'$,*

$$\Pr\left(\left\|p\,\hat{\Sigma}-\frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i\boldsymbol{x}_i^T\right\|\le\varepsilon\right)\ge 1-Cne^{-c\varepsilon^2 n}. \tag{11}$$

The strategy of the proof for Theorem 3.2 is as follows. According to Lemma 2.1, a scaled Tyler's M-estimator is a linear combination of $\boldsymbol{x}_i\boldsymbol{x}_i^T$, i.e., it can be written as $\sum_{i=1}^n \hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T$ (up to a scaling). Then Lemma 3.1 shows that $n\hat{w}_i$ converges to 1 uniformly, and based on the following matrix analysis, Theorem 3.2 can be concluded.

Proof of Theorem 3.2. We first prove that for $\varepsilon<c'$,

$$\Pr\left(\left\|\sum_{i=1}^n \hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T-\frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i\boldsymbol{x}_i^T\right\|\le\varepsilon\right)\ge 1-Cne^{-c\varepsilon^2 n}. \tag{12}$$

Let $\boldsymbol{B}_n=\sum_{i=1}^n(\hat{w}_i-\frac{1}{n})\boldsymbol{x}_i\boldsymbol{x}_i^T=\sum_{i=1}^n \hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T-\frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i\boldsymbol{x}_i^T$, then

$$\begin{aligned}\|\boldsymbol{B}_n\|&=\sup_{\|\boldsymbol{v}\|=1}\boldsymbol{v}^T\boldsymbol{B}_n\boldsymbol{v}=\sup_{\|\boldsymbol{v}\|=1}\sum_{i=1}^n(\hat{w}_i-\frac{1}{n})(\boldsymbol{v}^T\boldsymbol{x}_i)^2\\&\le\sup_{\|\boldsymbol{v}\|=1}\sum_{i=1}^n\left\|\hat{\boldsymbol{w}}-\frac{1}{n}\mathbf{1}\right\|_\infty(\boldsymbol{v}^T\boldsymbol{x}_i)^2\le\|n\hat{\boldsymbol{w}}-\mathbf{1}\|_\infty\|\frac{1}{n}\sum_{i=1}^n\boldsymbol{x}_i\boldsymbol{x}_i^T\|.\end{aligned}$$

Since $\|n\hat{\boldsymbol{w}}-\mathbf{1}\|_\infty\to 0$ with probability estimated in (10), and $\|\frac{1}{n}\sum_{i=1}^n\boldsymbol{x}_i\boldsymbol{x}_i^T\|$ is bounded above by $(1+2\sqrt{y})^2$ with probability $1-C\exp(-cn)$ [9, Theorem II.13], (12) is proved.

Second, since $\|\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T\|\le\left\|\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T-\frac{1}{n}\sum_{i=1}^n\boldsymbol{x}_i\boldsymbol{x}_i^T\right\|+\left\|\frac{1}{n}\sum_{i=1}^n\boldsymbol{x}_i\boldsymbol{x}_i^T\right\|$,

$$\Pr\left(\|\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T\|<C'\right)>1-Cn\exp(-cn). \tag{13}$$

Besides, $\mathrm{tr}(\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T)=\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i^T\boldsymbol{x}_i\to p$ in the same rate as in (13): applying the concentration of high-dimensional Gaussian measure on the sphere [4, Corollary 2.3], we have

$$\begin{aligned}&\max\left(\Pr\left(\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i^T\boldsymbol{x}_i<p(1-\varepsilon)\right),\Pr\left(\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i^T\boldsymbol{x}_i>p/(1-\varepsilon)\right)\right) \qquad (14)\\&\le\max\left(\Pr\left(\min_{1\le i\le n}\|\boldsymbol{x}_i\|^2<p(1-\varepsilon)\right),\Pr\left(\max_{1\le i\le n}\|\boldsymbol{x}_i\|^2>p/(1-\varepsilon)\right)\right)<ne^{-\varepsilon^2p/4}.\end{aligned}$$

Combining (13), (14) and (6),

$$\left\|\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T-p\,\hat{\Sigma}\right\|=\left\|\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T\right\|\left(1-p/\mathrm{tr}(\sum_{i=1}^n\hat{w}_i\boldsymbol{x}_i\boldsymbol{x}_i^T)\right) \tag{15}$$

converges in the same rate as specified in (12). (11) is then proved by combining (12), (15) and the triangle inequality. □

From the probabilistic estimation (11) we obtain a convergence rate of $O(\sqrt{\log n/n})$. In simulations we observe a rate of $O(1/\sqrt{n})$, which means our estimation might be off by a factor of $\sqrt{\log n}$.

3.1.2 Maronna's M-estimator

In this section we first state our assumptions for $u(x)$ in (4):

A1. $u : [0, \infty) \to (0, \infty)$ is nonnegative, $\psi(x) = xu(x)$ is increasing and $\lim_{x\to\infty} \psi(x) > y$.

A2. $u(x)$ is twice differentiable, and $xu'(x) < u(x)$.

We require assumption **A1** to ensure the existence and uniqueness of Maronna's M-estimator so that Lemma 2.2 can be applied. When $u(x)$ is nonnegative and $\psi(x) = xu(x)$ is increasing, the uniqueness condition in Section 2, i.e., $\rho(x)$ is non-decreasing and $\rho(e^x)$ is convex, are guaranteed (recall $\rho'(x) = nu(x)/2$). And the condition $\lim_{x\to\infty} \psi(x) > y$ guarantees the existence of Maronna's M-estimator as $p, n \to \infty$, as discussed in Section 2.

We require assumption **A2** for some technical steps in our proof, though we conjecture that our results about Maronna's M-estimator in this paper will still hold without this assumption.

Here we compare our assumption of $u(x)$ with the assumption in [7, 8]. Since $\big(Z, u(x)\big)$ in [7, 8] is equivalent to $\big(p\bar{\Sigma}, yu(x)\big)$ in our setting, their assumptions of $u(x)$ can be translated to:

- $u(x)$ is nonnegative, continuous and increasing.
- $\psi(x)$ is increasing and bounded, and $\frac{1}{y} < \lim_{x\to\infty} \psi(x) < \frac{1}{y^2}$.

There are three main differences between the assumptions of $u(x)$, and our assumptions allow some $u(x)$ that was not covered in their work. First, our assumption of $\lim_{x\to\infty} \psi(x)$ is less restrictive and allows it to be infinity. As a consequence, our theory allows some commonly used $u(x)$ such as $u(x) = x^\beta$ (see [27]). Second, our assumption is less restrictive in the sense that we replaced the assumption "$u(x)$ is nonincreasing" (i.e., $u'(x) \le 0$) by $xu'(x) < u(x)$. However, our assumption on the twice differentiability of $u(x)$ is more restrictive.

Based on these assumptions, we obtain the convergence of Maronna's M-estimator to a scaled version of $\boldsymbol{S}_n$ in operator norm.

Lemma 3.3. *If $\boldsymbol{x}_i \sim N(\boldsymbol{0}, \mathbf{I})$ for all $1 \le i \le n$, let $\psi(x) = xu(x)$, then there exists a solution $\{\bar{w}_i\}_{i=1}^n$ to (8) $\max_{1\le i\le n} |\bar{w}_i - \psi^{-1}(\frac{1}{y})|$ converges to 0 almost surely as $p, n \to \infty$. In particular, there exists $C, c, c' > 0$ such that for any $\varepsilon < c'$,*

$$\Pr\left(\max_{1\le i\le n} |\bar{w}_i - \psi^{-1}(1/y)| \le \varepsilon\right) \ge 1 - Cne^{-c\varepsilon^2 n}. \tag{16}$$

Theorem 3.4. *Suppose that $\{\boldsymbol{x}_i\}_{i=1}^n$ are i.i.d. sampled from $N(\mathbf{0}, \mathbf{I})$, $p, n \to \infty$ and $p/n \to y$, where $0 < y < 1$, and $\boldsymbol{x}_i \sim N(\mathbf{0}, \mathbf{I})$ for all $1 \le i \le n$, then a scaled Maronna's M-estimator converges to $\boldsymbol{S}_n$ in operator norm almost surely, and there exist $C, c, c' > 0$ such that for any $\varepsilon < c'$,*

$$\Pr\left(\left\|\frac{1}{n\,\psi^{-1}(1/y)}\,\bar{\Sigma} - \frac{1}{n}\sum_{i=1}^{n}\boldsymbol{x}_i\boldsymbol{x}_i^T\right\| \le \varepsilon\right) \ge 1 - Cne^{-c\varepsilon^2 n}. \tag{17}$$

3.1.3 Convergence of the largest eigenvalue of Tylers and Maronnas M-estimators

It follows from Theorems 3.2 and 3.4 that both $p\,\hat{\Sigma}$ and $\frac{1}{n\,\psi^{-1}(1/y)}\,\bar{\Sigma}$ converge to the sample covariance asymptotically in operator norm. Combining it with the convergence of the largest eigenvalue of the sample covariance [1, Theorem 5.11] to $(1+\sqrt{y})^2$ and the triangle inequality, we obtain the convergence of the largest eigenvalues of Tyler's and Maronna's M-estimators.

Corollary 3.5. *Under the settings of Theorems 3.2 and 3.4, we have a.s.*

$$\lim_{n\to\infty} p\|\hat{\Sigma}\| = \lim_{n\to\infty} \frac{1}{n\,\psi^{-1}(1/y)}\,\|\bar{\Sigma}\| = (1+\sqrt{y})^2.$$

3.2 More General Distributions

3.2.1 Tyler's M-estimator

In this section, we extend Theorem 3.2 from the setting of the normal distribution $N(\mathbf{0}, \mathbf{I})$ to elliptical distributions. We say that μ_p is an elliptical distribution, if μ_p can be characterized by $\mu_p(\boldsymbol{x}) = C(g_p)\det(\boldsymbol{T}_p)^{-1/2} g_p(\boldsymbol{x}^T\boldsymbol{T}_p^{-1}\boldsymbol{x})$, where $\boldsymbol{T}_p$ is a positive definite matrix in $\mathbb{R}^{p\times p}$, $g_p : [0,\infty)$ to $[0,\infty)$ satisfies $\int_0^\infty g_p(x)x^{p-1} < \infty$, and $C(g_p)$ is a normalization parameter that only depends on g_p.

When $\boldsymbol{T}_p$ is a scalar matrix, the distribution is isotropic and we call μ_p spherically symmetric distribution.

Our analysis is based on Theorem 3.2 and two properties of Tyler's M-estimator: 1. Tyler's M-estimator is invariant to the scaling of data set, i.e., if $\{\boldsymbol{x}_i\}_{i=1}^n$ are replaced by $\{c_i\boldsymbol{x}_i\}_{i=1}^n$ and $\{c_i\}_{i=1}^n$ are arbitrary numbers in $\mathbb{R}$, then $\hat{\Sigma}$ remains the same. 2. For any non-singular linear operator $\boldsymbol{T} : \mathbb{R}^p \to \mathbb{R}^p$, if Tyler's M-estimator for $\{\boldsymbol{x}_i\}_{i=1}^n$ is $\hat{\Sigma}$, then Tyler's M-estimator for $\{\boldsymbol{T}\boldsymbol{x}_i\}_{i=1}^n$ is $\boldsymbol{T}\hat{\Sigma}\boldsymbol{T}^T/\mathrm{tr}(\boldsymbol{T}\hat{\Sigma}\boldsymbol{T}^T)$.

Both properties can be obtained by verifying (2). To prove the first propriety, note that the LHS of (2) is unchanged if $\{\boldsymbol{x}_i\}_{i=1}^n$ is replaced by $\{c_i\boldsymbol{x}_i\}_{i=1}^n$. To prove the second property, one can show that (2) still holds when $\{\boldsymbol{x}_i\}_{i=1}^n$ and $\hat{\Sigma}$ are replaced by $\{\boldsymbol{T}\boldsymbol{x}_i\}_{i=1}^n$ and $\boldsymbol{T}\hat{\Sigma}\boldsymbol{T}/\mathrm{tr}(\boldsymbol{T}\hat{\Sigma}\boldsymbol{T})$.

Theorem 3.6. *If $\{\boldsymbol{x}_i\}_{i=1}^n$ are i.i.d. sampled from elliptical distribution $\mu_p(\boldsymbol{x}) = C(g_p)\det(\boldsymbol{T}_p)^{-1/2} g_p(\boldsymbol{x}^T\boldsymbol{T}_p^{-1}\boldsymbol{x})$, then we have the following property for Tyler's*

M-estimator: there exist $c, C, c' > 0$ such that for any $\varepsilon < c'$,

$$\Pr\left(\left\| p\boldsymbol{T}_p^{-1/2}\hat{\Sigma}\boldsymbol{T}_p^{-1/2}/\mathrm{tr}(\boldsymbol{T}_p^{-1/2}\hat{\Sigma}\boldsymbol{T}_p^{-1/2}) - \frac{p}{n}\sum_{i=1}^{n}\boldsymbol{y}_i\boldsymbol{y}_i^T \right\| \le \varepsilon\right) \ge 1 - Cne^{-c\varepsilon^2 n} \tag{18}$$

for $\boldsymbol{y}_i = \boldsymbol{T}_p^{-1/2}\boldsymbol{x}_i/\|\boldsymbol{T}_p^{-1/2}\boldsymbol{x}_i\|$.

Proof. First, since $\boldsymbol{x}_i$ are i.i.d. sampled from an elliptical distribution with covariance matrix $\boldsymbol{T}_p$, $\boldsymbol{y}_i = \boldsymbol{T}_p^{-1/2}\boldsymbol{x}_i/\|\boldsymbol{T}_p^{-1/2}\boldsymbol{x}_i\|$ are uniformly distributed over the $p-1$-dimensional unit sphere.

If we consider $\boldsymbol{y}_i$ as the projections of $\boldsymbol{x}_i \sim N(\boldsymbol{0}, \mathbf{I})$, the concentration of $N(\boldsymbol{0}, \mathbf{I})$ on the sphere with radius $\sqrt{p}$ [4, Corollary 2.3] and the boundedness of $\|\frac{1}{n}\sum_{i=1}^{n}\boldsymbol{x}_i\boldsymbol{x}_i^T\|$ from above [9, Theorem II.13] gives that for $\varepsilon < c'$,

$$\Pr\left(\left\| \frac{p}{n}\sum_{i=1}^{n}\boldsymbol{y}_i\boldsymbol{y}_i^T - \frac{1}{n}\sum_{i=1}^{n}\boldsymbol{x}_i\boldsymbol{x}_i^T \right\| < \varepsilon\right) <\geq 1 - Cne^{-c\varepsilon^2 n}. \tag{19}$$

Assuming the Tyler's M-estimator for $\{\boldsymbol{y}_i\}_{i=1}^n$ is $\hat{\Sigma}_y$, then Theorem 3.2 and (19) gives

$$\Pr\left(\left\| p\,\hat{\Sigma}_y - \frac{p}{n}\sum_{i=1}^{n}\boldsymbol{y}_i\boldsymbol{y}_i^T \right\| \le \varepsilon\right) \ge 1 - Cne^{-c\varepsilon^2 n}. \tag{20}$$

Applying Property 1 (scale invariance) of Tyler's M-estimator, $\hat{\Sigma}_y$ is also the Tyler's M-estimator for the set $\{\boldsymbol{T}_p^{-1/2}\boldsymbol{x}_i\}_{i=1}^n$. Applying Property 2,

$$\hat{\Sigma}_y = \boldsymbol{T}_p^{-1/2}\hat{\Sigma}\boldsymbol{T}_p^{-1/2}/\mathrm{tr}(\boldsymbol{T}_p^{-1/2}\hat{\Sigma}\boldsymbol{T}_p^{-1/2}). \tag{21}$$

Combining (20) and (21), Theorem 3.6 is proved. □

3.2.2 Maronna's M-estimator

In this section, we extend Theorem 3.4 from the setting of the normal distribution $N(\boldsymbol{0}, \mathbf{I})$ to non-isotropic Gaussian distributions. The model is more restrictive than the model of Tyler's M-estimator, since Maronna's M-estimator lacks Property 1 (scale invariance) of Tyler's M-estimator. We extend Theorem 3.4 to non-isotropic Gaussian distributions by applying a similar property to the Property 2 of Tyler's M-estimator: For any non-singular linear operator $\boldsymbol{T} : \mathbb{R}^p \to \mathbb{R}^p$, if Maronna's M-estimator for $\{\boldsymbol{x}_i\}_{i=1}^n$ is $\bar{\Sigma}$, then Tyler's M-estimator for $\{\boldsymbol{T}\boldsymbol{x}_i\}_{i=1}^n$ is $\boldsymbol{T}\hat{\Sigma}\boldsymbol{T}$.

Corollary 3.7. *If $\{\boldsymbol{x}_i\}_{i=1}^n$ are i.i.d. sampled from $N(\boldsymbol{0}, \boldsymbol{T}_p)$, where $\boldsymbol{T}_p$ is a positive definite matrix in $\mathbb{R}^{p\times p}$,. Then we have the following property for Tyler's M-estimator: there exist $c, C, c' > 0$ such that for any $\varepsilon < c'$,*

$$\Pr\left(\left\| \boldsymbol{T}_p^{-1/2}\Big(\frac{1}{n\,\psi^{-1}(1/y)}\,\bar{\Sigma} - \frac{1}{n}\sum_{i=1}^{n}\boldsymbol{x}_i\boldsymbol{x}_i^T\Big)\boldsymbol{T}_p^{-1/2} \right\| \le \varepsilon\right) \ge 1 - Cne^{-c\varepsilon^2 n}. \tag{22}$$

The distributions for data samples in this section can be compared to the model given in [7, Section II] and [8, Assumption 2]. For the simplicity of the discussion we only discuss [8, Assumption 2], which assumes that $\boldsymbol{x}_i \in \mathbb{R}^p$ is defined by $\sqrt{\tau_i}\boldsymbol{A}_N\boldsymbol{y}_i$, where $\boldsymbol{y}_i$ has independent entries with zero mean and unit variance, and τ_i follows from some distribution.

While [8] covers more models than Corollary 3.7, we note that our proof only depends Lemma 5.2. That is, our proof can be applied to any distribution that satisfies Lemma 5.2. Since the distribution in [8] satisfies [8, Lemma 6], which is equivalent to Lemma 5.2, our proof can also be applied to their models.

3.3 Empirical Spectral Density

3.3.1 Tyler's M-estimator

This section investigates the distribution of the eigenvalues of Tyler's M-estimator, i.e., its empirical spectral density. We follow the setting of previous sections and present two corollaries, where the first corollary proves the conjecture proposed in [13] that the empirical spectral density converges to the Marčenko-Pastur distribution when $\{\boldsymbol{x}_i\}_{i=1}^n$ are drawn from $N(\mathbf{0}, \mathbf{I})$, and the second corollary gives the limiting distribution under the setting of elliptical distributions.

Corollary 3.8. *If $\{\boldsymbol{x}_i\}_{i=1}^n$ are i.i.d. sampled from spherically symmetric distributions $C(g_p)g_p(\|\boldsymbol{x}\|^2)$, then the empirical spectral density of $p\hat{\Sigma}$ converges to the Marčenko-Pastur distribution.*

To visualize Corollary 3.8, we simulated the case $n = 20000$ and $p = 4000$ with Gaussian distribution $N(\mathbf{0}, \mathbf{I})$, and Figure 1 shows that the empirical spectral density of $p\hat{\Sigma}$ is well approximated by the corresponding Marčenko-Pastur distribution.

Lemma 3.9. *Assume a set of matrices $\{\boldsymbol{A}_n\}_{n\geq 1}$ with size $k_n \times k_n$, and with empirical spectral density converging to a continuous distribution ρ, and another sequence of matrices $\{\boldsymbol{B}_n\}_{n\geq 1}$ such that $\boldsymbol{B}_n$ is also of size $k_n \times k_n$ and $\|\boldsymbol{B}_n\| \to 0$. Then the empirical spectral density of $\{\boldsymbol{A}_n + \boldsymbol{B}_n\}_{n\geq 1}$ also converges to ρ.*

Proof of Corollary 3.8. The proof follows from Theorem 3.2 and Lemma 3.9, and the proof of Lemma 3.9 will be given later in Section 5.

First, due to Property 1 in Section 3.2, it suffices to consider the case $\boldsymbol{x}_i \sim N(\mathbf{0}, \mathbf{I})$. Then Corollary 3.8 is proved by combining Theorem 3.2, Lemma 3.9 and the fact that the empirical density of $\frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i\boldsymbol{x}_i^T$ converges to (3) as $p, n \to \infty$ [18]. □

Next we extend the analysis to general elliptical distributions.

Corollary 3.10. *Suppose $\{\boldsymbol{x}_i\}_{i=1}^n$ are i.i.d. sampled from elliptical distribution $C(g_p)g_p(\boldsymbol{x}^T\boldsymbol{T}_p^{-1}\boldsymbol{x})$, and the empirical spectral density of $\boldsymbol{T}_p$ converges to H.*

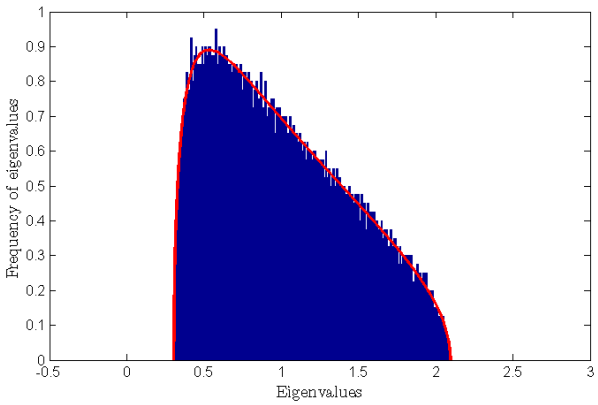

Figure 1: The empirical spectral density of a properly scaled Tyler's M-estimator when $n = 20000$ and $p = 4000$, and $\{\boldsymbol{x}_i\}_{i=1}^n$ are i.i.d. drawn from $N(\mathbf{0}, \mathbf{I})$. The red line represents the Marčenko-Pastur distribution for $y = 1/5$.

Then the empirical spectral density of $\mathrm{tr}(\boldsymbol{T}_p)\hat{\Sigma}$ *converges to* ρ*, whose Stieltjes transform* $s(z)$ *satisfies*

$$s(z) = \int \frac{1}{t(1 - y - yz\, s(z)) - z}\, \mathrm{d}H(t). \tag{23}$$

Proof. Let

$$\boldsymbol{B}_p = p\boldsymbol{T}_p^{-1/2}\hat{\Sigma}\boldsymbol{T}_p^{-1/2}/\mathrm{tr}(\boldsymbol{T}_p^{-1/2}\hat{\Sigma}\boldsymbol{T}_p^{-1/2}) - \frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T, \tag{24}$$

where $\boldsymbol{z}_i = h_i \cdot \boldsymbol{y}_i = h_i \cdot \boldsymbol{T}_p^{-1/2}\boldsymbol{x}_i/\|\boldsymbol{T}_p^{-1/2}\boldsymbol{x}_i\|$, and $h_i \sim \sqrt{\chi_p^2}$. Then Theorem 3.6 and the convergence of h_i to $\sqrt{p}$ imply

$$\hat{\Sigma} = \frac{\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \boldsymbol{B}_p)\boldsymbol{T}_p^{1/2}}{\mathrm{tr}\left(\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \boldsymbol{B}_p)\boldsymbol{T}_p^{1/2}\right)},$$

where $\|\boldsymbol{B}_p\| \to 0$.

Since $\|\boldsymbol{B}_p\| \to 0$ and $\boldsymbol{z}_i$ are i.i.d from $N(\mathbf{0}, \mathbf{I})$, $\mathrm{tr}\left(\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \boldsymbol{B}_p)\boldsymbol{T}_p^{1/2}\right) \to \mathrm{tr}(\boldsymbol{T}_p)$ almost surely. Therefore we only need to prove that the empirical spectral density of $\mathrm{tr}\left(\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \boldsymbol{B}_p)\boldsymbol{T}_p^{1/2}\right)\hat{\Sigma} = \boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \boldsymbol{B}_p)\boldsymbol{T}_p^{1/2}$ converges to ρ.

Since $\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \boldsymbol{B}_p)\boldsymbol{T}_p^{1/2} - \boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T - \|\boldsymbol{B}_p\|\mathbf{I})\boldsymbol{T}_p^{1/2}$ is positive definite, the eigenvalues of $\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \boldsymbol{B}_p)\boldsymbol{T}_p^{1/2}$ is bounded below by $\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T - \|\boldsymbol{B}_p\|\mathbf{I})\boldsymbol{T}_p^{1/2}$; similarly it is bounded above by

$\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \|\boldsymbol{B}_p\|\mathbf{I})\boldsymbol{T}_p^{1/2}$. Combining it with the fact that the eigenvalues of $\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T$ are almost surely bounded by $[1-\sqrt{y}-t, 1+\sqrt{y}+t]$ for any $t > 0$ [24, Corollary 5.35], the eigenvalues of $\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T + \boldsymbol{B}_p)\boldsymbol{T}_p^{1/2}$ are bounded below and above by $\boldsymbol{T}_p^{1/2}(1 - \frac{\|\boldsymbol{B}_p\|}{1-\sqrt{y}-t})(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T)\boldsymbol{T}_p^{1/2}$ and $\boldsymbol{T}_p^{1/2}(1 + \frac{\|\boldsymbol{B}_p\|}{1-\sqrt{y}-t})(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T)\boldsymbol{T}_p^{1/2}$ almost surely. By the convergence of the ESD of $\boldsymbol{T}_p^{1/2}(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\boldsymbol{z}_i^T)\boldsymbol{T}_p^{1/2}$ to ρ [1, (6.1.2)] and the convergence of $\|\boldsymbol{B}_p\|$ to 0, Corollary 3.10 is proved.

□

3.3.2 Maronna's M-estimator

This section investigates the distribution of the eigenvalues of Maronna's M-estimator, when data is sampled from a Gaussian distribution. The analysis follows from the proof of Theorem 3.6 and Corollary 3.7.

Corollary 3.11. *Suppose $\{\boldsymbol{x}_i\}_{i=1}^n$ are i.i.d. sampled from Gaussian distribution $N(\mathbf{0}, \boldsymbol{T}_p)$, where $\boldsymbol{T}_p$ is a positive definite matrix in $\mathbb{R}^{p\times p}$, and the empirical spectral density of $\boldsymbol{T}_p$ converges to H. Then the empirical spectral density of $\frac{1}{n\,\psi^{-1}(1/y)}\bar{\Sigma}$ converges to ρ, whose Stieltjes transform $s(z)$ satisfies*

$$s(z) = \int \frac{1}{t(1-y-yz\,s(z)) - z}\,\mathrm{d}H(t). \tag{25}$$

In particular, if $\boldsymbol{T}_p = \mathbf{I}$ for all p, $\frac{1}{n\,\psi^{-1}(1/y)}\bar{\Sigma}$ converges to the Marčenko-Pastur distribution in (3).

3.4 Simulations on the Spike Covariance Model

In this section we present results of numerical simulations. The purpose is to illustrate that in the spike covariance model, Tyler's M-estimator outperforms the sample covariance when the data sets are generated from an elliptical distribution with a heavy tail. The setup of the two experiments is as follows:

1. $\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_{200}$ are i.i.d. sampled from a Multivariate t-distribution $t_1(\mathbf{0}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma} \in \mathbb{R}^{20\times 20}$ and $\boldsymbol{\Sigma} = \mathrm{diag}(2, 1, 1, \cdots, 1)$.

2. $\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_{200}$ are i.i.d. sampled from a Multivariate t-distribution $t_1(\mu, \mathbf{I})$, where $\mathbf{I}$ is a matrix of size 20×20 and $\mu = (1, 0, 0, \cdots, 0) \in \mathbb{R}^{20}$. In this example, we assume that the signal $(1, 0, 0, \cdots, 0)$ is contaminated by a noise from Multivariate t-distribution $t_1(\mathbf{0}, \mathbf{I})$.

For both cases, there is one spike and the direction of the "spike" is $(1, 0, \cdots, 0)$. To examine the performance of the sample covariance estimator and Tyler's M-estimator, we plot the distribution of their eigenvalues in Figure 2. The figure indicates that the eigenvalues of Tyler's M-estimator has one significant spike and the other eigenvalues are in the 'bulk'. However, the number of spikes is unclear from the eigenvalues of the sample covariance estimator.

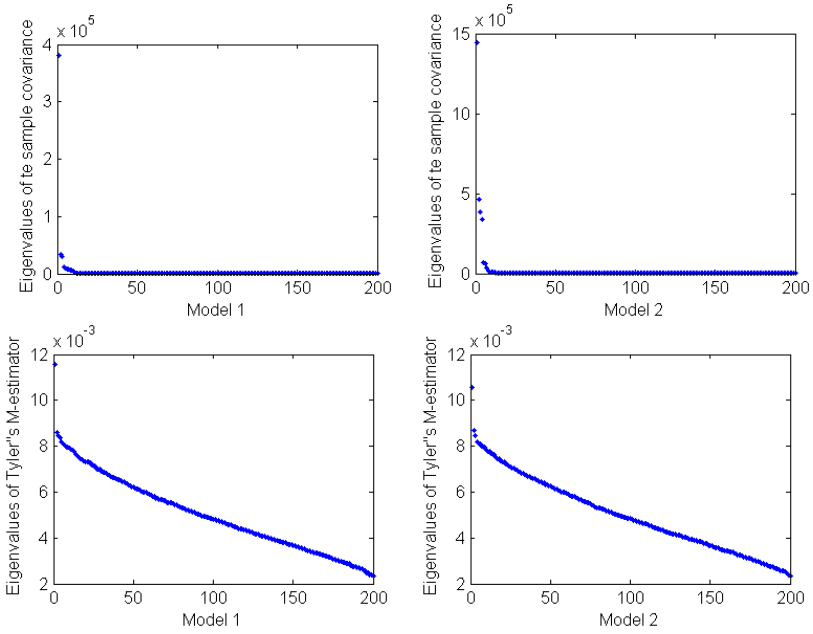

Figure 2: The empirical distributions of the eigenvalues of Tyler's M-estimators and the sample covariance estimators in simulated models proposed in Section 3.4.

We also record the mean correlations between the top eigenvectors of the covariance estimators and the underlying spike (as well as the standard deviations) over 100 runs in Table 1. The table shows that the sample covariance failed to recover the spike in the sense that the top eigenvector has a small correlation with the true spike. In comparison, the top eigenvector of Tyler's M-estimator is consistently close to the true spike. In summary, these simulations suggest the superiority of Tyler's M-estimator over the sample covariance when data samples are generated from a heavy-tailed elliptical distribution.

	Model 1	Model 2
Tyler's M-estimator	0.903(0.013)	0.903(0.035)
Sample covariance	0.080(0.057)	0.057(0.012)

Table 1: Mean and standard deviation (in parenthesis) of the correlation between the leading eigenvector of the population covariance matrix and that of the estimated covariance matrix.

4 Summary

We established that Maronna's M-estimator and Tyler's M-estimator converge in operator norm to $\boldsymbol{S}_n$ as $p, n \to \infty$ and $p/n \to y$, $0 < y < 1$, where data samples follow the distribution of $N(\mathbf{0}, \mathbf{I})$. We also extended the result to elliptical distribution for Tyler's M-estimator and non-isotropic Gaussian distributions for Maronna's M-estimator, and proved the conjecture that the empirical spectral density of Tyler's M-estimator converges to the Marčenko-Pasture distribution.

There are several possible future directions of this work. First, it might be possible to strengthen the results. For example, we would like to know if a more careful analysis can prove the convergence of Maronna's estimator without the assumption **A2**. Second, in simulations we observe the rate of convergence of the M-estimator to $\boldsymbol{S}_n$ is $1/\sqrt{n}$, while the current theoretical analysis only gives the order of $O(\sqrt{\log n/n})$, and we would like to find an approach that gives the better empirical convergence rate. At last, it would be interesting to theoretically quantify the behavior of Tyler's and Marrona's M-estimators in the spiked covariance model, which includes the analysis of the distribution of the top eigenvalue for the null cases and the analysis of the nonnull case.

5 Proof of Lemmas

5.1 Proof of Lemma 2.1

We start with the definition

$$(\hat{z}_1, \hat{z}_2, \cdots, \hat{z}_n) = \underset{\sum_{i=1}^n z_i = 1}{\arg\min} \log\det(\sum_{i=1}^n e^{z_i} \boldsymbol{x}_i \boldsymbol{x}_i^T) \tag{26}$$

and

$$\hat{\Sigma}_z = \sum_{i=1}^{n} e^{\hat{z}_i} \boldsymbol{x}_i \boldsymbol{x}_i^T. \tag{27}$$

The solution to (26) is unique, which follows from the convexity of the objective function (see [25, Lemma 4]). Besides, noticing the equivalence between (26) and (7) (by plugging $w_i = e^{z_i}/(\sum_{i=1}^n e^{z_i})$ and $z_i = \log w_i - (\sum_{i=1}^n \log w_i - 1)/n$), there exists $c_1 > 0$ such that $\hat{\Sigma}_z = c_1 \hat{\Sigma}$.

Next we will prove that $\hat{\Sigma}_z$ satisfies

$$\sum_{i=1}^{n} \frac{\boldsymbol{x}_i \boldsymbol{x}_i^T}{\boldsymbol{x}_i^T \hat{\Sigma}_z^{-1} \boldsymbol{x}_i} = c\hat{\Sigma}_z, \text{ for some } c > 0. \tag{28}$$

By checking the directional derivative of the objective function in (26), for any $(\delta_1, \delta_2, \cdots, \delta_n)$ with $\sum_{i=1}^n \delta_i = 0$,

$$\sum_{i=1}^{n} \delta_i e^{\hat{z}_i} \boldsymbol{x}_i^T \hat{\Sigma}_z^{-1} \boldsymbol{x}_i = 0.$$

Therefore, there exists c_2 such that

$$e^{\hat{z}_i} \boldsymbol{x}_i^T \hat{\Sigma}_z^{-1} \boldsymbol{x}_i = c_2, \text{ for all } 1 \le i \le n. \tag{29}$$

Therefore, (28) is proved by applying (29) and (27):

$$\sum_{i=1}^{n} \frac{\boldsymbol{x}_i \boldsymbol{x}_i^T}{\boldsymbol{x}_i^T \hat{\Sigma}_z^{-1} \boldsymbol{x}_i} = \sum_{i=1}^{n} e^{\hat{z}_i} \boldsymbol{x}_i \boldsymbol{x}_i^T / c_2 = \hat{\Sigma}_z / c_2,$$

Since $\hat{\Sigma}_z = c_1 \hat{\Sigma}$, (28) also holds when $\hat{\Sigma}_z$ is replaced by $\hat{\Sigma}$:

$$\sum_{i=1}^{n} \frac{\boldsymbol{x}_i \boldsymbol{x}_i^T}{\boldsymbol{x}_i^T \hat{\Sigma}^{-1} \boldsymbol{x}_i} = c\hat{\Sigma}, \text{ for some } c > 0. \tag{30}$$

At last, we will prove that $\hat{\Sigma}$ satisfies the definition of Tyler's M-estimator in (2), that is, the constant c in (30) is given by $c = n/p$. For the objective function

$$F(\Sigma) = \sum_{i=1}^{n} \log(\boldsymbol{x}_i^T \Sigma^{-1} \boldsymbol{x}_i) + c \log\det(\Sigma),$$

its derivative with respect to Σ^{-1} is given by

$$\sum_{i=1}^{n} \boldsymbol{x}_i^T (\boldsymbol{x}_i^T \Sigma^{-1} \boldsymbol{x}_i)^{-1} \boldsymbol{x}_i - c\Sigma.$$

Therefore, $\hat{\Sigma}$ is a stationary point of $F(\Sigma)$. Since $F(\Sigma)$ is geodesically convex (argument follows directly from [25, 26]), $\hat{\Sigma}$ is the global minimizer of $F(\Sigma)$.

However, the minimizer of $F(\Sigma)$ exists only when $c = n/p$. Since $F(a\mathbf{I}) = \sum_{i=1}^n \log(\boldsymbol{x}_i^T \boldsymbol{x}_i) - n \log a + c\, p \log a$, we have

$$F(a\mathbf{I}) \to -\infty \quad \begin{cases} \text{as } a \to 0, \text{ if } c > n/p \\ \text{as } a \to \infty, \text{ if } c < n/p \end{cases}.$$

Therefore, the constant c in (30) is given by $c = n/p$, and Lemma 2.1 is proved.

5.2 Proof of Lemma 3.1

We start with an outline of the proof, which consists of three parts. First, we rewrite the constrained optimization problem (7) to the problem of finding the root of $g(\boldsymbol{w})$, which will be defined in (31). Since the root of $g(\boldsymbol{w})$ is $n\hat{\boldsymbol{w}} - 1$, we only need to show the convergence of the root of $g(\boldsymbol{w})$. Second, we will show that $g(\mathbf{0})$ converges to $\mathbf{0}$, $\nabla g(\mathbf{0})$ is large and the variation of $\nabla g(\boldsymbol{w})$ is bounded. Finally, we will use a perturbation analysis and the observations on $g(\mathbf{0})$ and $\nabla g(\boldsymbol{w})$ to show that the root of $g(\boldsymbol{w})$ converges to $\mathbf{0}$.

The proof depends on Lemma 5.2, Lemma 5.3 and Lemma 5.1, and their proofs are postponed to subsequent sections.

Lemma 5.1. *For a function $f(\boldsymbol{w}) : \mathbb{R}^p \to \mathbb{R}^p$, assume that $\nabla f(\mathbf{0}) = \mathbf{I}$, and $\|\nabla f(\boldsymbol{w}) - \nabla f(\mathbf{0})\|_\infty = \max_{i \le i \le p} \|\nabla f_i(\boldsymbol{w}) - \nabla f_i(\mathbf{0})\|_\infty < C_5 \|\boldsymbol{w}\|_\infty$ for $\|\boldsymbol{w}\|_\infty \le 1$, and $\|f(\mathbf{0})\|_\infty < \min(1/9C_5, 1/3)$. Then there exists $\tilde{\boldsymbol{w}}$ such that $\|\tilde{\boldsymbol{w}}\|_\infty < 3\|f(\mathbf{0})\|_\infty$ and $f(\tilde{\boldsymbol{w}}) = \mathbf{0}$.*

Lemma 5.2. *If $\boldsymbol{x}_i \sim N(\mathbf{0}, \mathbf{I})$ for all $1 \le i \le n$, and $\boldsymbol{S} = \frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i \boldsymbol{x}_i^T$, then there exists $c, C, c' > 0$ such that for any $\varepsilon < c'$,*

$$\Pr\left(\max_{1 \le i \le n} |\frac{1}{p} \boldsymbol{x}_i^T \boldsymbol{S}^{-1} \boldsymbol{x}_i - 1| < \varepsilon \right) \ge 1 - Cne^{-c\varepsilon^2 n}.$$

Lemma 5.3. *For the $n \times n$ matrix $\boldsymbol{A}$ defined by $\boldsymbol{A}_{ij} = \frac{1}{n\,p}(\boldsymbol{x}_i^T \boldsymbol{S}^{-1} \boldsymbol{x}_j)^2$, (a) $\|\boldsymbol{A}\|_\infty < 2$ with probability $1 - Cn\exp(-cn)$.*

(b) There exists $c = c(p, n) > 0$ and $C_2 = C_2(y) > 0$ such that $\|(\mathbf{I} - \boldsymbol{A} + c\mathbf{1}\mathbf{1}^T)^{-1}\|_\infty < C_2$ with probability $1 - Cn\exp(-cn)$.

We start the first part of the proof with the construction of $g(\boldsymbol{w})$. We let

$$g(\boldsymbol{w}) = \nabla G(\boldsymbol{w} + \mathbf{1}), \tag{31}$$

where

$$G(\boldsymbol{w}) = -\sum_{i=1}^n \log w_i + \frac{n}{p} \log\det(\sum_{i=1}^n w_i \boldsymbol{x}_i \boldsymbol{x}_i^T) + \frac{c_0}{2}(\sum_{i=1}^n w_i - n)^2, \tag{32}$$

and the constant c_0 will be specified later before (48).

It is easy to prove that the minimizer of $G(\boldsymbol{w})$ and the zeros of $\nabla G(\boldsymbol{w})$ must satisfies $\sum_{i=1}^n w_i = n$ (otherwise $n\boldsymbol{w}/(\sum_{i=1}^n w_i)$ is a better minimizer and $\nabla G(\boldsymbol{w})$ is nonzero). Therefore minimizing (32) is equivalent to minimizing $-\sum_{i=1}^n \log w_i + \frac{n}{p}\log\det(\sum_{i=1}^n w_i \boldsymbol{x}_i\boldsymbol{x}_i^T)$ with constraint $\sum_{i=1}^n w_i = n$, which is the same as (7) except for the constraint. Noticing that a scaling of $\boldsymbol{w}$ increases $-\sum_{i=1}^n \log w_i + \frac{n}{p}\log\det(\sum_{i=1}^n w_i \boldsymbol{x}_i\boldsymbol{x}_i^T)$ by a constant only depending on the scale, the minimizer of (32) is unique and it is $n\hat{\boldsymbol{w}}$, where $\hat{\boldsymbol{w}}$ is defined in (7). By the convexity of its equivalent problem (26), the root of $g(\boldsymbol{w})$ is also unique and it is $n\hat{\boldsymbol{w}} - 1$.

For the second part of the proof, we start by proving that $g(\boldsymbol{0})$ is small. By calculation, the i-th component of function $g(\boldsymbol{w})$ is

$$g_i(\boldsymbol{w}) = -\frac{1}{w_i+1} + \frac{n}{p}\boldsymbol{x}_i^T(n\boldsymbol{S} + \sum_{i=1}^n w_i\boldsymbol{x}_i\boldsymbol{x}_i^T)^{-1}\boldsymbol{x}_i + c_0(\sum_{i=1}^n w_i).$$

Applying Lemma 5.2,

$$\Pr\left(\|g(\boldsymbol{0})\|_\infty < \varepsilon\right) \geq 1 - Cne^{-c\varepsilon^2 n}. \tag{33}$$

Now we will prove that $\nabla g(\boldsymbol{0})$ is bounded from below. By calculation, its (i,j)-th entry is

$$\left(\nabla g(\boldsymbol{w})\right)_{i,j} = I(i=j)\frac{1}{(w_i+1)^2} - \frac{n}{p}\Big(\boldsymbol{x}_i^T\big(n\boldsymbol{S} + \sum_{i=1}^n w_i\boldsymbol{x}_i\boldsymbol{x}_i^T\big)^{-1}\boldsymbol{x}_j\Big)^2 + c_0.$$

Applying Lemma 5.3,

$$\|(\nabla g(\boldsymbol{0}))^{-1}\|_\infty < C_2 \text{ with probability } 1 - Cne^{-cn}. \tag{34}$$

Now we bound the variation of $\nabla g(\boldsymbol{w})$ in the region $\|\boldsymbol{w}\|_\infty < 1/2$. Apply $|\frac{1}{(w_i+1)^2} - 1| < 3|w_i - 1| \leq 3\|\boldsymbol{w}\|_\infty$ and coordinatewise comparison,

$$|\nabla_{i,j} g(\boldsymbol{w}) - \nabla_{i,j} g(\boldsymbol{0})| \leq I(i=j)\left(3\|\boldsymbol{w}\|_\infty\right) + 3\|\boldsymbol{w}\|_\infty \cdot \frac{n}{p}|\boldsymbol{A}_{ij}|.$$

Therefore, the variation of $\nabla g(\boldsymbol{w})$ is bounded by

$$\|\nabla g(\boldsymbol{w}) - \nabla g(\boldsymbol{0})\|_\infty < (3 + 3n\|\boldsymbol{A}\|_\infty/p)\|\boldsymbol{w}\|_\infty. \tag{35}$$

At last we finish the third part of the proof of Lemma 3.1 by applying Lemma 5.1 to $f(\boldsymbol{w}) = (\nabla g(\boldsymbol{0}))^{-1} g(\boldsymbol{w}/2)$. It is easy to verify that $\nabla f(\boldsymbol{0}) = \mathbf{I}$. Due to (33) and (34), $\|f(\boldsymbol{0})\|_\infty \leq \|(\nabla g(\boldsymbol{0}))^{-1}\|_\infty \|g(\boldsymbol{0})\|_\infty \to 0$ in the same rate as in (33) and $\|f(\boldsymbol{0})\|_\infty < \min(1/9C_5, 1/3)$ holds with probability $1 - Cne^{-cn}$. Due to (33), (35), and the boundedness of $\|\boldsymbol{A}\|_\infty$ (Lemma 5.3), $\|\nabla f(\boldsymbol{w}) - \nabla(\boldsymbol{0})\|_\infty < C_5\|\boldsymbol{w}\|_\infty$ also holds with probability $1 - Cne^{-cn}$. Therefore the assumption in Lemma 5.1 holds with probability $1 - Cne^{-cn}$ and there exists $\tilde{\boldsymbol{w}}$ such that $f(\tilde{\boldsymbol{w}}) = 0$ and

$$\|\tilde{\boldsymbol{w}}\|_\infty < 3\|f(\boldsymbol{0})\|_\infty. \tag{36}$$

When $f(\tilde{\boldsymbol{w}}) = 0$, we have $g(2\tilde{\boldsymbol{w}}) = 0$ and by previous discussion $2\tilde{\boldsymbol{w}} = n\hat{\boldsymbol{w}} - 1$. therefore (36) gives

$$\|n\hat{\boldsymbol{w}} - 1\|_\infty < 6\|f(\mathbf{0})\|_\infty.$$

Since $\|f(\mathbf{0})\|_\infty$ converges to 0 in the rate as in (33), $\|n\hat{\boldsymbol{w}} - 1\|_\infty$ converges in the same rate and Lemma 3.1 is proved.

5.3 Proof of Lemma 3.3

Proof. Define $\bar{g}(\boldsymbol{w}) : \mathbb{R}^n \to \mathbb{R}^n$ by $\bar{g}_j(\boldsymbol{w}) = w_j - \boldsymbol{x}_j^T \big(\sum_{i=1}^n u(w_i)\boldsymbol{x}_i\boldsymbol{x}_i^T \big)^{-1} \boldsymbol{x}_j$ for all $1 \le j \le n$, then $\bar{\boldsymbol{w}} = (\bar{w}_1, \bar{w}_2, \cdots, \bar{w}_n)$ is a root of $\bar{g}(\boldsymbol{w})$. Let $\tilde{\boldsymbol{w}} = \psi^{-1}(1/y)\mathbf{1}$, we will prove that

1. $\Pr(\|\bar{g}(\tilde{\boldsymbol{w}})\|_\infty < \varepsilon) > 1 - Cne^{-c\varepsilon^2 n}$ for any $\varepsilon < c'$.
2. $\Pr(\|\big(\nabla\bar{g}(\tilde{\boldsymbol{w}})\big)^{-1}\|_\infty < C') > 1 - Cne^{-cn}$.
3. $\Pr(\|\nabla\bar{g}(\tilde{\boldsymbol{w}}) - \nabla\bar{g}(\boldsymbol{w})\|_\infty < C'\|\boldsymbol{w} - \tilde{\boldsymbol{w}}\|_\infty) > 1 - Cne^{-cn}$ for any $\|\boldsymbol{w} - \tilde{\boldsymbol{w}}\|_\infty < c'$.

The first point can be proved by applying Lemma 5.2. As for the second point, we have

$$\nabla_{ij}\bar{g}(\tilde{\boldsymbol{w}}) = \mathbf{I} + \frac{u'(\psi^{-1}(1/y))}{u(\psi^{-1}(1/y))^2}\big(\boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_j\big)^2.$$

Since $\sum_{j=1}^n \big(\boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_j\big)^2 = \boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_i$, and $\boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_i$ converges to $1/y$ with probability (Lemma 5.2), we have

$$\begin{aligned}
&\sum_{j=1}^n \frac{u'(\psi^{-1}(1/y))}{u(\psi^{-1}(1/y))^2}\big(\boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_j\big)^2 \to \frac{u'(\psi^{-1}(1/y))}{u(\psi^{-1}(1/y))^2} \cdot 1/y \\
=&\frac{u'(\psi^{-1}(1/y))}{u(\psi^{-1}(1/y))^2} \cdot u(\psi^{-1}(1/y))\psi^{-1}(1/y) = \frac{u'(\psi^{-1}(1/y))\,\psi^{-1}(1/y)}{u(\psi^{-1}(1/y))}.
\end{aligned}$$

Since $|u'(x)x/u(x)| < 1$, we have $\Big\| \frac{u'(\psi^{-1}(1/y))}{u(\psi^{-1}(1/y))^2}\big(\boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_j\big)^2 \Big\|_\infty < 1$ with exponential probability, and by

$$\begin{aligned}
&\Big\|\mathbf{I} + \frac{u'(\psi^{-1}(1/y))}{u(\psi^{-1}(1/y))^2}\big(\boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_j\big)^2\Big\|_\infty \le \sum_{i=0}^{\infty} \Big\| \frac{u'(\psi^{-1}(1/y))}{u(\psi^{-1}(1/y))^2}\big(\boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_j\big)^2\Big\|_\infty^i \\
=&\frac{1}{1 - \big\| \frac{u'(\psi^{-1}(1/y))}{u(\psi^{-1}(1/y))^2}\big(\boldsymbol{x}_i(n\boldsymbol{S})^{-1}\boldsymbol{x}_j\big)^2\big\|_\infty},
\end{aligned}$$

the second point is proved.

As for the third point, we note that the j-th component of the gradient of $\bar{g}_k$ is

$$\nabla_j\big(\bar{g}_k(\tilde{\boldsymbol{w}})\big) = \delta(j = k) + u'(w_j)\Big(\boldsymbol{x}_k^T\big(\sum_{i=1}^n u(w_i)\boldsymbol{x}_i\boldsymbol{x}_i^T\big)^{-1}\boldsymbol{x}_j\Big)^2,$$

u is twice differentiable, and $\|\boldsymbol{A}\|_\infty$ is bounded with high probability, therefore the third point holds.

Apply Lemma 5.1 with $f(\boldsymbol{w}) = \big(\nabla\bar{g}(\tilde{\boldsymbol{w}})\big)^{-1}\bar{g}(\boldsymbol{w})$, then Lemma 3.3 follows the same procedure as in the third part of the proof of Lemma 3.1. □

5.3.1 Proof of Lemma 5.1

Proof. When $\|\boldsymbol{w}\|_\infty \le 1$,

$$f_j(\boldsymbol{w}) - f_j(\boldsymbol{0}) = \int_{t=0}^{1} \left\langle \boldsymbol{e}_j \boldsymbol{w}^T, \nabla f(t\,\boldsymbol{w}) \right\rangle \mathrm{d}t \tag{37}$$

$$= \int_{t=0}^{1} \left\langle \boldsymbol{e}_j \boldsymbol{w}^T, \nabla f(t\,\boldsymbol{w}) - \nabla f(\boldsymbol{0}) + \mathbf{I} \right\rangle \mathrm{d}t = w_j + \int_{t=0}^{1} \boldsymbol{w}^T (\nabla f(t\,\boldsymbol{w}) - \nabla f(\boldsymbol{0})) \boldsymbol{e}_j \,\mathrm{d}t$$

$$\le w_j + \| \int_{t=0}^{1} \boldsymbol{w}^T (\nabla f(t\,\boldsymbol{w}) - \nabla f(\boldsymbol{0}))\|_\infty \le w_j + C_5 \|\boldsymbol{w}\|_\infty^2.$$

Similarly

$$f_j(\boldsymbol{w}) - f_j(\boldsymbol{0}) \ge -C_5 \|\boldsymbol{w}\|_\infty^2 + w_j. \tag{38}$$

To prove it, we consider the continuous mapping $h(\boldsymbol{w}) = \boldsymbol{w} - f(\boldsymbol{w})/(4+9C_5)$ and will prove that h maps $\mathcal{A}$ to itself, where

$$\mathcal{A} = \{\boldsymbol{w} : \boldsymbol{w} \in [-3\eta, 3\eta]^n\} \text{ and } \eta = \|f(\boldsymbol{0})\|_\infty.$$

1. $|w_i| < 2\eta$. Then apply (37) and (38) (they are applicable since for any $\boldsymbol{w} \in \mathcal{A}$, $\|\boldsymbol{w}\|_\infty \le 1$), we have $|f_i(\boldsymbol{w})| < |f_i(\boldsymbol{0})| + C_5\|\boldsymbol{w}\|_\infty^2 + |w_i| \le \eta + C_5(3\eta)^2 + 3\eta < (4+9C_5)\eta$ ($\eta^2 < \eta$ since $\eta < 1$). Therefore, $|h_i(\boldsymbol{w})| \le |w_i| + |f_i(\boldsymbol{w})|/(4+9C_5) \le 3\eta$.

2. $w_i > 2\eta$, then applying (38),

$$f_i(\boldsymbol{w}) \ge -|f_i(\boldsymbol{0})| + w_i - C_5\|\boldsymbol{w}\|_\infty^2 \ge -\eta + 2\eta - C_5(3\eta)^2.$$

Since $\eta < 1/9C_5$, we have $f_i(\boldsymbol{w}) < 0$ and therefore $h_i(\boldsymbol{w}) \le w_i \le 3\eta$.

Similar to case 1 we can prove that $h_i(\boldsymbol{w}) \ge -3\eta$. Therefore $|h_i(\boldsymbol{w})| < 3\eta$.

3. Similar to case 2, when $w_i < -2\eta$, $|h_i(\boldsymbol{w})| < 3\eta$.

Therefore the continuous mapping h maps the convex, compact set $\mathcal{A}$ to itself. By Schauder fixed point theorem, $h(\boldsymbol{x})$ has a fixed point in $\mathcal{A}$ and Lemma 5.1 is proved with $\tilde{\boldsymbol{w}}$ being the fixed point. □

5.3.2 Proof of Lemma 5.2

Assuming the SVD decomposition of $\boldsymbol{X}$ is $\boldsymbol{X} = \boldsymbol{U}\Sigma\boldsymbol{V}^T$, where $\boldsymbol{U} \in \mathbb{R}^{n\times p}$ and $\boldsymbol{U}^T\boldsymbol{U} = \mathbf{I}$. Since $\boldsymbol{x}_i \sim N(\boldsymbol{0}, \mathbf{I})$ for all $1 \le i \le n$, $\boldsymbol{U}$ is uniformly distributed over the space of all orthogonal $n \times p$ matrices. Since

$$\boldsymbol{X}\boldsymbol{S}^{-1}\boldsymbol{X} = (\boldsymbol{U}\Sigma\boldsymbol{V}^T)(\frac{1}{n}\boldsymbol{V}\Sigma^2\boldsymbol{V}^T)^{-1}(\boldsymbol{U}\Sigma\boldsymbol{V}^T), \tag{39}$$

if we write the row of $\boldsymbol{U}$ by $\boldsymbol{u}_1, \boldsymbol{u}_2, \cdots, \boldsymbol{u}_n$, then $\frac{1}{n}\boldsymbol{x}_i\boldsymbol{S}^{-1}\boldsymbol{x}_i = \boldsymbol{u}_i^T\boldsymbol{u}_i = \|\boldsymbol{u}_i\|^2$.

Since $\boldsymbol{U}$ can be considered as the first p columns of a random $n\times n$ orthogonal matrix (with haar measure over the set of all $n \times n$ orthogonal matrices), $\boldsymbol{u}_i$ can be considered as the first p entries from a random vector of length n that is sampled from the uniform sphere in $\mathbb{R}^n$.

Therefore, $\|\boldsymbol{u}_i\|^2 \sim \sum_{j=1}^p g_j^2 / \sum_{j=1}^n g_j^2$ for i.i.d. random variables $\{g_j\}_{j=1}^n \sim N(0,1)$. Applying [4, Corollary 2.3], we have

$$\Pr\left(\sum_{i=1}^n g_i^2 \geq \frac{n}{1-\varepsilon}\right) \leq e^{-\varepsilon^2 n/4} \tag{40}$$

and

$$\Pr\left(\sum_{i=1}^n g_i^2 \leq n(1-\varepsilon)\right) \leq e^{-\varepsilon^2 n/4}, \tag{41}$$

therefore

$$\begin{aligned}
&\Pr\left(\frac{p(1-\varepsilon)^2}{n} \leq \|\boldsymbol{u}_1\|^2 \leq \frac{p}{n(1-\varepsilon)^2}\right) \geq \Pr\left(p(1-\varepsilon) \leq \sum_{i=1}^p g_i^2 \leq \frac{p}{1-\varepsilon}\right) \\
&+ \Pr\left(n(1-\varepsilon) \leq \sum_{i=1}^n g_i^2 \leq \frac{n}{1-\varepsilon}\right) \geq 1 - 2e^{-\varepsilon^2 p/4} - 2e^{-\varepsilon^2 n/4}.
\end{aligned}$$

For $\varepsilon \leq 0.1$, we have

$$\begin{aligned}
&\Pr\left(\max_{1\leq i\leq n} |\frac{1}{p}\boldsymbol{x}_i^T \boldsymbol{S}^{-1}\boldsymbol{x}_i - 1| \leq \varepsilon\right) \geq 1 - n\Pr\left(|\|\boldsymbol{u}_1\|^2 - \frac{p}{n}| > \frac{p}{n}\varepsilon\right) \\
\geq& 1 - n\left(1 - \Pr\left(\frac{p(1-\varepsilon/3)^2}{n} \leq \|\boldsymbol{u}_1\|^2 \leq \frac{p}{n(1-\varepsilon/3)^2}\right)\right) \qquad (42) \\
\geq& 1 - 2ne^{-\varepsilon^2 p/36} - 2ne^{-\varepsilon^2 n/36}, \qquad (43)
\end{aligned}$$

where the second inequality follows from $1 - 3\varepsilon \leq (1-\varepsilon)^2$ and $\frac{1}{(1-\varepsilon)^2} \leq 1 + 3\varepsilon$.

5.3.3 Proof of Lemma 5.3

(a) Since $\|\boldsymbol{A}\|_\infty = \max_{1\leq i\leq n}(\sum_{1\leq j\leq n} \boldsymbol{A}_{ij})$, and

$$\sum_{1\leq j\leq n} \boldsymbol{A}_{ij} = \sum_{1\leq j\leq n} \frac{1}{np}\boldsymbol{x}_i^T \boldsymbol{S}^{-1}\boldsymbol{x}_j\boldsymbol{x}_j^T \boldsymbol{S}^{-1}\boldsymbol{x}_i = \boldsymbol{x}_i^T \boldsymbol{S}^{-1}(\sum_{1\leq j\leq n} \boldsymbol{x}_j\boldsymbol{x}_j^T)\boldsymbol{S}^{-1}\boldsymbol{x}_i/np \tag{44}$$

$$= \boldsymbol{x}_i^T \boldsymbol{S}^{-1}(n\boldsymbol{S})\boldsymbol{S}^{-1}\boldsymbol{x}_i/np = \boldsymbol{x}_i^T \boldsymbol{S}^{-1}\boldsymbol{x}_i/p, \tag{45}$$

it follows from (43) with $\varepsilon = 0.1$ that $\|\boldsymbol{A}\|_\infty < 2$ holds with probability $1 - Cn\exp(-cn)$.

(b) We first prove that there exists $C_3 = C_3(y)$ such that

$$\|\boldsymbol{A} - c_0\mathbf{1}\mathbf{1}^T\|_\infty \leq C_3 < 1 \text{ with probability } 1 - Cn\exp(-cn). \tag{46}$$

We start with the proof of (46) with another lemma:

Lemma 5.4. *There exists a $c_4 > 0$ such that with probability $1 - C\exp(-cn)$,*

$$\sum_{j=1}^{n} I(\boldsymbol{x}_1^T \boldsymbol{x}_j > c_4\sqrt{p}) > 0.75n.$$

There exists $C_4 = C_4(y)$ such that $\|\boldsymbol{S}\| < C_4$ with probability $1-Cn\exp(-cn)$ [9, Theorem II.13]. Therefore $\boldsymbol{x}_i^T \boldsymbol{S}^{-1}\boldsymbol{x}_j \geq \boldsymbol{x}_i^T \boldsymbol{x}_j/C_4$ and Lemma 5.4 implies that for any $1 \leq i \leq n$:

$$\sum_{j=1}^{n} I(\boldsymbol{x}_i^T \boldsymbol{S}^{-1}\boldsymbol{x}_j > c_4\sqrt{p}/C_4) > 0.75 \text{ with probability } 1 - C\exp(-cn). \quad (47)$$

Let $c_0 = (c_4/C_4)^2/n$, then (47) implies

$$\sum_{1\leq j\leq n} |\boldsymbol{A}_{i,j} - c| \leq \sum_{1\leq j\leq n} |\boldsymbol{A}_{i,j}| - 0.25c\,n \leq \boldsymbol{x}_i^T \boldsymbol{S}^{-1}\boldsymbol{x}_i/p - 0.25(c_4/C_4)^2, \quad (48)$$

where the last step follows from (45).

Applying the estimation of $\boldsymbol{x}_i^T \boldsymbol{S}^{-1}\boldsymbol{x}_i/p$ in (43) and a union bound argument over all $1 \leq i \leq n$ to (48), (46) is proved for $C_3 = 1 + \eta - 0.25(c_4/C_4)^2$.

Lemma 5.3(b) follows from (46) with $C_2 = \frac{1}{1-C_3}$, where the expansion of $(\mathbf{I} - \boldsymbol{A} + c\mathbf{1}\mathbf{1}^T)^{-1}$ is valid since $\|\boldsymbol{A} + c\mathbf{1}\mathbf{1}^T\| \leq \|\boldsymbol{A} + c\mathbf{1}\mathbf{1}^T\|_\infty < 1$:

$$\|(\mathbf{I} - \boldsymbol{A} + c\mathbf{1}\mathbf{1}^T)^{-1}\|_\infty \leq \sum_{i=0}^{\infty} \|\boldsymbol{A} - c\mathbf{1}\mathbf{1}^T\|_\infty^i = \sum_{i=0}^{\infty} C_3^i = \frac{1}{1-C_3}. \quad (49)$$

5.3.4 Proof of Lemma 5.4

We first show that there exists c_4 such that for all p,

$$E(I(|\boldsymbol{x}_1^T \boldsymbol{x}_2| > c_4\sqrt{p})) \geq 0.85. \quad (50)$$

WLOG we rotate $\boldsymbol{x}_1$ such that it is nonzero only at the first coordinate, and $\boldsymbol{x}_2 = (g_1, g_2, ..., g_p)$ where $g_i \sim N(0,1)$. Then $|\boldsymbol{x}_1^T \boldsymbol{x}_2| = |g_1|\,\|\boldsymbol{x}_1\|$.

Notice that $\|\boldsymbol{x}_1\|^2$ is the sum of p independent χ_1^2 distribution and $E\chi_1^2 = 1$, by central limit theorem, $\|\boldsymbol{x}_1\| \leq \sqrt{2p}$ with probability $1 - Ce^{-cn}$. Besides, $\Pr(|g_1| > \sqrt{2}\,c_4) \geq 0.85$ for $c_4 = \Phi^{-1}(1 - 0.85/2)/\sqrt{2}$. Therefore (50) is proved by combining the estimations on $|g_1|$, $\boldsymbol{x}_1$ and $|\boldsymbol{x}_1^T \boldsymbol{x}_2| = |g_1|\,\|\boldsymbol{x}_1\|$.

To obtain Lemma 5.4 from (50), we apply Hoeffding's inequality to the indicator function $I(|\boldsymbol{x}_i^T \boldsymbol{x}_j| > c_4\sqrt{p})$ over all $1 \leq j \leq n, j \neq i$.

5.4 Proof of Lemma 3.9

Denoting the k-th eigenvalue of any matrix $\boldsymbol{A}$ by $\lambda_k(\boldsymbol{A})$, then [5, Corollary III.4.2] gives

$$\lambda_k(\boldsymbol{A}_n + \boldsymbol{B}_n) - \lambda_k(\boldsymbol{A}_n) \leq \|\boldsymbol{B}_n\|. \quad (51)$$

Assuming the empirical spectral density of $\boldsymbol{A}_n$ and $\boldsymbol{A}_n + \boldsymbol{B}_n$ are ρ_n and ρ'_n, then (51) implies

$$\int_a^b \rho'_n(x)\,\mathrm{d}x \leq \int_{a-\|\boldsymbol{B}_n\|}^{b+\|\boldsymbol{B}_n\|} \rho_n(x)\,\mathrm{d}x.$$

Since $\|\boldsymbol{B}_n\| \to 0$, for any $\varepsilon > 0$,

$$\lim_{n\to\infty} \sup \int_a^b \rho'_n(x)\,\mathrm{d}x \leq \int_{a-\varepsilon}^{b+\varepsilon} \rho(x)\,\mathrm{d}x.$$

By the continuity of ρ, $\lim_{n\to\infty} \sup \int_a^b \rho'_n(x)\,\mathrm{d}x \leq \int_a^b \rho(x)$. Similarly we can prove that $\lim_{n\to\infty} \inf \int_a^b \rho'_n(x)\,\mathrm{d}x \geq \int_a^b \rho(x)$, and therefore $\lim_{n\to\infty} \int_a^b \rho'_n(x)\,\mathrm{d}x = \int_a^b \rho(x)$ and Lemma 3.9 is proved.

Acknowledgement

A. Singer was partially supported by Award Number FA9550-12-1-0317 and FA9550-13-1-0076 from AFOSR, by Award Number R01GM090200 from the NIGMS, and by Award Number LTR DTD 06-05-2012 from the Simons Foundation.

References

[1] Bai, Z., & Silverstein, J.W. 2009. *Spectral Analysis of Large Dimensional Random Matrices*. Springer series in statistics. Springer.

[2] Baik, Jinho, & Silverstein, Jack W. 2006. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of Multivariate Analysis*, **97**(6), 1382 – 1408.

[3] Baik, Jinho, Ben Arous, Gérard, & Péché, Sandrine. 2005. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, **33**(5), 1643–1697.

[4] Barvinok, Alexander. 2005. *Math 710: Measure Concentration*. Lecture notes, Department of Mathematics, University of Michigan.

[5] Bhatia, R. 1997. *Matrix Analysis*. Graduate Texts in Mathematics. Springer Verlag.

[6] Chen, Yilun, Wiesel, A., & Hero, A.O. 2011. Robust Shrinkage Estimation of High-Dimensional Covariance Matrices. *Signal Processing, IEEE Transactions on*, **59**(9), 4097 –4107.

[7] Couillet, Romain, Pascal, Frédéric, & Silverstein, Jack W. 2012. Robust M-Estimation for Array Processing: A Random Matrix Approach. *CoRR*, **abs/1204.5320**.

[8] Couillet, Romain, Pascal, Frédéric, & Silverstein, Jack W. 2013. The Random Matrix Regime of Maronna's M-estimator with elliptically distributed samples. *CoRR*, **abs/1311.7034**.

[9] Davidson, Kenneth R., & Szarek, Stanislaw J. 2001. Local operator theory, random matrices and Banach spaces. *Pages 317–366 of: Handbook of the geometry of Banach spaces, Vol. I.* Amsterdam: North-Holland.

[10] Dümbgen, Lutz. 1998. On Tyler's M-Functional of Scatter in High Dimension. *Annals of the Institute of Statistical Mathematics*, **50**(3), 471–491.

[11] El Karoui, Noureddine. 2009. Concentration of measure and spectra of random matrices: Applications to correlation matrices, elliptical distributions and beyond. *The Annals of Applied Probability*, **19**(6), 2362–2405.

[12] Frahm, Gabriel, & Glombek, Konstantin. 2012. Semicircle law of Tyler's M-estimator for scatter. *Statistics & Probability Letters*, **82**(5), 959 – 964.

[13] Frahm, Gabriel, & Jaekel, Uwe. 2007. *Tyler's M-estimator, random matrix theory, and generalized elliptical distributions with applications to finance.* Discussion Papers in Statistics and Econometrics 2/07. University of Cologne, Department for Economic and Social Statistics.

[14] Hardt, M., & Moitra, A. 2012. Algorithms and Hardness for Robust Subspace Recovery. *arXiv preprint arXiv:1211.1041*.

[15] Johnstone, Iain M. 2001. On the Distribution of the Largest Eigenvalue in Principal Components Analysis. *The Annals of Statistics*, **29**(2), pp. 295–327.

[16] Kent, John T., & Tyler, David E. 1991. Redescending M-Estimates of Multivariate Location and Scatter. *The Annals of Statistics*, **19**(4), pp. 2102–2119.

[17] Maronna, R. A. 1976. Robust M-Estimators of Multivariate Location and Scatter. *The Annals of Statistics*, **4**(1), pp. 51–67.

[18] Marčenko, V A, & Pastur, L A. 1967. DISTRIBUTION OF EIGENVALUES FOR SOME SETS OF RANDOM MATRICES. *Mathematics of the USSR-Sbornik*, **1**(4), 457.

[19] Nadler, Boaz. 2008. Finite sample approximation results for principal component analysis: A matrix perturbation approach. *The Annals of Statistics*, **36**(6), 2791–2817.

[20] Ollila, E., & Koivunen, V. 2003. Robust antenna array processing using M-estimators of pseudo-covariance. *Pages 2659–2663 vol.3 of: Personal, Indoor and Mobile Radio Communications, 2003. PIMRC 2003. 14th IEEE Proceedings on*, vol. 3.

[21] Ollila, E., & Tyler, D.E. 2012 (June). Distribution-free detection under complex elliptically symmetric clutter distribution. *Pages 413–416 of: Sensor Array and Multichannel Signal Processing Workshop (SAM), 2012 IEEE 7th.*

[22] Soshnikov, Alexander. 2002. A Note on Universality of the Distribution of the Largest Eigenvalues in Certain Sample Covariance Matrices. *Journal of Statistical Physics*, **108**(5-6), 1033–1056.

[23] Tyler, David E. 1987. A Distribution-Free M-Estimator of Multivariate Scatter. *The Annals of Statistics*, **15**(1), pp. 234–251.

[24] Vershynin, Roman. 2011. *Introduction to the non-asymptotic analysis of random matrices.* Available at http://arxiv.org/abs/1011.3027v7.

[25] Wiesel, A. 2012. Geodesic Convexity and Covariance Estimation. *Signal Processing, IEEE Transactions on*, **60**(12), 6182 –6189.

[26] Zhang, T. 2012. Robust subspace recovery by geodesically convex optimization. *arXiv preprint arXiv:1206.1386.*

[27] Zhang, Teng, Wiesel, A., & Greco, M.S. 2013. Multivariate Generalized Gaussian Distribution: Convexity and Graphical Models. *Signal Processing, IEEE Transactions on*, **61**(16), 4141–4148.